

**DIAGNOSING FAULTS AND ERRORS FROM A DATA REPOSITORY USING  
DIRECTED GRAPHS**

**FIELD OF THE INVENTION**

The present invention is directed to fault diagnosis from a data repository and, more particularly, to encoding subsets of possible questions and answers of the repository as directed graphs, and when the repository is remote from the user, downloading the subsets to substantially reduce the number of transmissions between the user and the remote repository.

**BACKGROUND DESCRIPTION**

Fault diagnosis can be accomplished by several techniques, including decision trees, case based analysis, and Bayesian networks. Each of these techniques requires narrowing the set of possible causes and solutions by answering a series of questions. When the information is stored in a remote data repository, considerable time is spent transmitting questions and subsequent answers between the user and the data repository.

A fault diagnosis session consists of the user responding to a series of questions (or sets of questions). The response to each question (or set of questions) reduces the set of possible causes of the fault and determines the next question (or set of questions) to be posed. This process repeats until one of the following conditions is true:

the number of possible causes of the fault is reduced to one (in which case, the cause is uniquely determined);

the number of possible causes of the fault is reduced to zero (in which

case, the cause cannot be determined); or

no further questions are available that will further reduce the number of causes (in which case, there remain multiple possible causes of the fault).

Each cycle of the diagnosis session has the following three steps:

(1) Select a question (or set of questions) to present to the user and determine the set of acceptable user responses. Note that the specific algorithm used to select one question (or set of questions) from the many potential questions distinguishes the various approaches to the problem known as case based reasoning, Bayesian networks, decision trees, etc. The specific approach is not important to the subject invention; any question selection algorithm may be employed. Similarly unimportant to the subject invention is the particular weighting scheme or measure employed to evaluate each potential question's suitability to be posed next.

(2) Present the selected question (or set of questions) to the user along with a description of the set of acceptable responses. Note that depending on the number of acceptable responses, various presentation methods may be employed to prompt the user's response. If the number is small, the system might display an itemized list of all acceptable responses as tokens or icons for the user to select. If the list is large, the system might display a text entry field for free-form entry of the response along with a description of the set of acceptable responses. With free-form entry, a parser and validation function must establish that the user's response is valid. The present invention extends to all methods of describing the valid set of responses and all methods of presenting them to the user for selection.

1 (3) Collect the user's response.

2 The basic unit of communication between the system and user in step (2) of the  
3 diagnosis session is the question to be posed and the set of potential valid answers to the  
4 question. The basic unit of communication between the user and system in step (3) of the  
5 diagnosis session is the user's actual response.

6 The question selection process in step (1) may require, as input, a list of all  
7 previously posed questions along with the actual user response to each question. These  
8 question-answer pairs form the current state of the process.

9 The question selection process in step (1) also requires as input:

10 a set of possible causes of faults;

11 a set of questions that can be posed to the user to diagnose the fault; and

12 a set of valid answers for each question.

13 In the most trivial scenario, one pass through the process, selecting and presenting  
14 one question (and establishing one question-answer pair), will suffice to resolve the fault.  
15 However, most scenarios will require at least one more pass through the process with a  
16 follow-up question (establishing an additional question-answer pair) prior to resolving the  
17 fault. In the general scenario, this process will repeat many times, presenting a new  
18 question-answer set to the user each time (and establishing an additional question-answer  
19 pair each time) before reaching one of the terminal conditions listed above.

20 When the database, the question selection algorithm, and the presentation method  
21 all reside on one dedicated single-user machine, latency due to data flows between the  
22 system and user are minimized. There are no network dependencies to slow the session  
23 (or cause it to completely fail) as data moves between the system and user. Similarly,  
24 state information is directly available to all steps of the process. The drawback to the  
25 single-user solution is the potentially large database of faults, questions, and answers that

1 must be downloaded to the stand-alone client. The latency due to downloading and  
2 installing a large database may more than offset the gains of local data flows.

3 In contrast to the single-client solution, the standard client-server solution places  
4 the database and all three steps of the fault diagnosis process upon the server. It offloads  
5 to the client only the extremely simple sub-tasks of displaying the question and potential  
6 answers (all pre-formatted on the server) and capturing the user's selection. Due to the  
7 simplistic use of the client, client-server solutions do not generally allow users to enter  
8 free-form text since this requires a sophisticated set of processing functions (parser,  
9 validation, and error processing functions) to ensure the user enters valid text. Instead, all  
10 valid responses must be enumerated explicitly by the server. This can be awkward when  
11 the set of valid responses is large (e.g., all valid twelve digit error codes).

12 Since the database and all three steps of the fault diagnosis process reside on the  
13 server, standard client-server solutions transmit one question (and its accompanying set of  
14 valid responses) from the server to the client during each pass through the process. The  
15 user's choice is transmitted back to the server, which then repeats the process, selecting a  
16 new question and sending it again to the client. Thus, each pass through the process  
17 requires two network transmissions and one server interaction.

18 This results in  $2n$  network transmissions and  $n$  server processes for  $n$   
19 question-answer interactions. The many server interactions slow server response and  
20 impede scalability to multiple users. The many network transmissions further slow  
21 response to the user, in addition to increasing the number of opportunities for failure in  
22 the middle of the session. Additionally, the process state information must either be  
23 transmitted with each interaction (further slowing response) or be stored on the server  
24 (further impeding server scalability). Storing the process state information on the server  
25 also means that this information must eventually be cleaned up, creating time-out  
26 problems for clients who have taken a break in the middle of a fault diagnosis session.

## 27 SUMMARY OF THE INVENTION

1           It is therefore an object of the present invention to provide a fault diagnosis  
2   technique which reduces latency due to downloading and installing a large database and  
3   also reduces the number of transmissions between user and the remote repository.  
4   According to the present invention, the necessity for transmitting each question and  
5   answer is reduced or eliminated and, by encoding the questions and answers, the latency  
6   is reduced. The present invention places all three steps of the process on the client,  
7   gaining all of the advantages of the single-client solution. However, it avoids the  
8   single-client solution drawback of the potentially large database download by  
9   downloading portions of the database onto the client as needed. These portions are  
10   significant enough in size to allow the process to run through many passes without  
11   returning to the server, but small enough to avoid delays in the diagnosis session due to  
12   the network latency. This substantially reduces the number of network transmissions and  
13   server interactions needed during a diagnosis session. It results in substantially better  
14   performance for the user and much better scalability for the server.

15           The technique comprises encoding the complete set of symptoms and diagnoses of  
16   the system as a graph, subdividing this graph into multiple subgraphs, and then loading  
17   subgraphs only as needed to continue with a diagnosis session.

18           An example embodiment of this invention uses Java classes to store the graphs of  
19   the repository, load them on demand, and download them when the repository is remote.  
20   The object oriented nature of Java supports the aggregation and encoding of commonly  
21   occurring sequences of question-answer pairs and related symptoms and diagnoses into  
22   one Java class. Java's ability to dynamically load additional object code at run-time as  
23   Java classes (a quality known as "late-binding") supports the incremental loading of  
24   additional graphs as needed. And when the repository is remote, Java's ability to  
25   remotely download a class (or package of classes) in one transmission minimizes network  
26   interactions needed to retrieve subsequent question-answer pairs.

27           Another example embodiment of this invention uses the dynamically  
28   downloadable quality of Java to dynamically move the question selection algorithm to the  
29   client from the server, offloading the processing from the server and obviating the return  
30   trip to the server for its processing ability. Similarly, the presentation functions are

1 moved to the client, allowing sophisticated client-side validation of free-form text entry  
2 with no server interaction.

3 Another example embodiment uses code generated by any other programming  
4 language, especially if the programming language is object oriented and supports late  
5 binding and on-demand downloading of classes as does Java.

## 6 BRIEF DESCRIPTION OF THE DRAWINGS

7 The foregoing and other objects, aspects and advantages will be better understood  
8 from the following detailed description of a preferred embodiment of the invention with  
9 reference to the drawings, in which:

10 Fig. 1 is a block diagram of a repository of directed graphs which taken together  
11 encode the symptoms and diagnoses of the fault diagnosis system of the subject  
12 invention,

13 Fig. 2 is a flow diagram showing the logic of the fault diagnosis system of the  
14 subject invention employing the repository of Figure 1 during a fault diagnosis session,  
15 and

16 Fig. 3 is an example of a block diagram of placement of system components in  
17 another embodiment of the subject invention when the repository is remote and directed  
18 graphs are downloaded as needed to a local agent running the logic of the fault diagnosis  
19 system.

## 20 DESCRIPTION OF THE INVENTION

21 The present invention provides methods, systems and apparatus for fault diagnosis  
22 which reduces latency due to downloading and installing a large database and reduces the  
23 number of transmissions between user and the remote repository. The necessity for  
24 transmitting each question and answer is reduced or eliminated and, by encoding the

1 questions and answers, the latency is reduced. The present invention places all the steps  
2 of the process upon the client, gaining all of the advantages of the single-client solution.  
3 However, it avoids the single-client solution drawback of the potentially large database  
4 download by downloading portions of the database onto the client as needed. These  
5 portions are significant enough in size to allow the process to run through many passes  
6 without returning to the server, but small enough to avoid delays in the diagnosis session  
7 due to the network latency. This substantially reduces the number of network  
8 transmissions and server interactions needed during a diagnosis session, and results in  
9 substantially better performance for the user and much better scalability for the server.

10 An example of a method comprises encoding the complete set of symptoms and  
11 diagnoses of the system as a graph, subdividing this graph into multiple subgraphs, and  
12 then loading subgraphs only as needed to continue with a diagnosis session.

13 A fault diagnosis session begins with the receipt of a description of the fault as a  
14 set of symptoms. This description may be used by any appropriate method (e.g., a  
15 method which indexes subgraphs by the symptoms they encode) to determine an  
16 appropriate subgraph to load with which to begin the diagnosis session. If no method has  
17 been defined or if the method fails to map the description to a unique subgraph, a  
18 previously designated default subgraph may be loaded to begin the session or the user  
19 may simply be prompted to choose a subgraph based on descriptions provided to the user  
20 of subgraphs which may serve as starting points.

21 If a diagnosis is immediately evident from the symptoms so far received, the  
22 currently loaded subgraph encodes this as a terminal node labeled with the diagnosis.  
23 Otherwise, further symptoms need to be received to diagnose the fault. The subgraph  
24 directs the probing for the next symptom, encoding the symptom as a node (representing  
25 a question that can be asked) and one or more edges leading from the node (each  
26 representing an acceptable answer to the question) to another node in the subgraph where  
27 the processing repeats as above until a diagnosis is found. An entirely new area of  
28 probing is encoded in the subgraph by a node linking from this subgraph to another  
29 subgraph, causing the new subgraph to be loaded and processed as just described.

30 Referring now to the drawings, and more particularly to Figure 1, there is shown a

1 block diagram of a repository of multiple graphs which encodes all of the known  
2 symptoms and faults of the target system being diagnosed. Database block 110 represents  
3 the complete repository of graphs and block 120 is the expansion of a typical graph of  
4 symptoms and faults in the repository. Although the method used to map symptoms and  
5 faults to particular graphs is not dictated by the subject invention, several organizational  
6 methods are typical: a graph might encode a subset of related symptoms that tend to  
7 co-occur in a diagnosis session or it might encode a set of related faults. Commonly, a  
8 subgraph will fully diagnose one particular subsystem of the overall faulty target system.

9 Each symptom in the graph comprises a question about the state of some attribute  
10 in the faulty system and an answer to the question representing the value of that attribute.  
11 Questions are encoded in the graph as nodes of type question and represented in Figure 1  
12 by diamond shaped nodes (e.g., nodes 130 and 131). Answers are encoded in the graph  
13 and represented in Figure 1 by labeled edges leaving a node of type *question* (e.g., edge  
14 140). A symptom is therefore encoded as a node of type question and an *answer* edge  
15 from the node (e.g., node 130 and edge 140 taken together encode a symptom). Each  
16 answer edge leads either to another node of type question (e.g., node 130 leads to node  
17 131 via edge 140), representing the next question to pose in the fault diagnosis session, or  
18 to one of several additional types of nodes:

19 a node of type *diagnosis* (hexagon shaped in Figure 1, e.g., node 150),  
20 representing one definitive diagnosis of the fault and optionally supplying an  
21 action plan to remedy the fault,

22 a node of type *call-graph* (round shaped in Figure 1, e.g., node 160),  
23 connecting one graph to another graph, allowing composition of graphs, or

24 a node of type *functional-branch* (rectangle shaped in Figure 1, e.g., node  
25 170), allowing a transfer of control to any other node in the current graph  
26 where the new node is the computed value of a function (associated with the  
27 functional-branch node) of the current state of the diagnosis session, where the



1 current state is represented by the set of all <question, answer> pairs formed  
2 from questions already answered in the current session.

3 Some embodiments may use additional node types to encode additional  
4 information useful in directing fault diagnosis operations. For example:

5 a node of type test performs a test (or directs the user to perform a test) on the  
6 faulty system to obtain an answer to a question,

7 a node of type lookup obtains an answer to a question by querying a source of  
8 information external to the diagnostic process and the user (e.g., a remote  
9 database), and

10 a node of type state-transformation modifies the current state by applying a  
11 function associated with the state-transformation node to the current state to  
12 modify the current state. This allows various forms of reasoning to be added  
13 to the diagnosis system, e.g., adding additional <question, answer> pairs to the  
14 current state based on the presence of symptoms already received is a form of  
15 automatic deduction of new symptoms.

16 Each of these additional node types would have non-labeled edges connecting  
17 them directly to the next node in the graph to be processed. Some embodiments may  
18 therefore choose to encode the information included in them directly as part of the next  
19 node rather than adding additional node types to the graph.

20 While the answer edges leading from a node of type question determine a set of  
21 enumerable answers to the question, some questions may not allow an easily enumerable  
22 set of answers, either because the set of answers is continuous or it is discrete but very  
23 large. To accommodate this situation, in some embodiments, an edge may be designated  
24 as *free-form* (e.g., edge 180), allowing any value to be the question's answer. Some  
25 embodiments may further allow an optional *validation* function to be associated with a

1 free-form edge to restrict the values to a set computable by the validation function.

2 One node of type question in each graph is designated the *root node* (node 130 in  
3 the diagram), creating a rooted and directed graph of symptoms and diagnoses. Typically  
4 (but not required by this invention), one graph that connects to all other graphs (even if  
5 only indirectly via other graphs), using multiple call-graph nodes, will be designated the  
6 *root graph*, thereby forming one large rooted directed graph of subgraphs capable of  
7 diagnosing all known faults of the target system. If no graph is designated as the root,  
8 then an auxiliary index (keyed on the root question node of at least one graph) might be  
9 used to access a dynamically selected root graph with which to begin a diagnosis session.

10 Referring now to Figure 2, there is shown a flow diagram of the logic of the  
11 process of fault diagnosis using the repository of directed graphs of Figure 1 according to  
12 the invention. The input to the process is the repository of all directed graphs diagnosing  
13 faults of a system.

14 In the following description, the following definitions are used:

15  $R$  is the repository of all directed graphs diagnosing faults of a system.

16  $G$  is the current directed graph being processed.

17  $N$  is the current node being processed.

18  $Q$  is the current question being posed.

19  $A$  is the answer to the current question  $Q$ .

20  $S$  is the current state of the diagnosis session, a set of pairs of the form  
21  $\langle \text{question}, \text{answer} \rangle$ , where there is one pair for each question already  
22 answered in the session.

1           The process begins in function block 210 by receiving a set of symptoms as an  
2   initial description of the fault being diagnosed, setting the current state  $S$  to be those  
3   symptoms, identifying one graph from the repository  $R$  which will serve as the root graph  
4   and initializing current graph  $G$  to be that graph. As described above, any method may be  
5   used to identify a root graph, e.g., indexing the graphs by initially received symptoms or  
6   simply designating one graph from which all other graphs can be reached. Graph  $G$  is  
7   then retrieved from the repository in function block 220. In function block 230,  $N$  is  
8   initialized to be the root node of graph  $G$ . The node type of  $N$  is identified in decision  
9   block 240. If  $N$  is of type *diagnosis*, then the process returns the contents of the node as  
10   the diagnosis of the fault. If the current node type is not of type diagnosis then one  
11   performs a particular node type operation. The particular node type operation is  
12   dependent upon the node type identified in the step of identifying the node type of the  
13   current node. This is repeated until the node type of the current node is of type diagnosis.

14           If  $N$  is of type *call-graph*, then, in function block 250,  $G$  is set to be the graph  
15   associated with the call-graph node and processing continues iteratively with this new  
16   graph at function block 220. If  $N$  is of type *functional-branch*, then, in block 260, the  
17   function associated with the functional-branch node is evaluated over the current state  $S$   
18   of the diagnosis session and the value of the function is assigned to  $N$ . The overall  
19   process continues iteratively with this new node at decision block 240. Otherwise,  $N$  is of  
20   type *question*, in which case  $Q$  is set to be the current question which is posed in function  
21   block 270,  $A$  is set to the answer to the question which is collected in function block 280,  
22   and state  $S$  of the session is updated with a pair of the form  $\langle Q, A \rangle$ . In addition to  
23   directly asking the user for the answer to the question, different embodiments of the  
24   invention might use some combination of additional node types (e.g., those described  
25   above: test, lookup, and state-transformation) to answer the question by performing a test  
26   on the faulty system, querying a source external to the faulty system, or deducing the  
27   answer from symptoms already received. In function block 290, the edge of the graph  
28   leaving  $N$  which is labeled with answer  $A$  is traversed to reach a new node in  $G$ ,  $N$  is  
29   initialized to be this new node, and processing continues iteratively with this new node at

1 decision block 220. If no node is labeled *A* but there is a *free-form* edge, that edge is  
2 traversed to reach the new node as above.

3 Referring now to Figure 3, there is shown a block diagram of the system  
4 components according to the invention when the repository of directed graphs is remote  
5 from the process running the logic of the fault diagnosis system. Database block 310 is  
6 the remote repository *R* including the complete set of directed graphs which taken  
7 together encode the symptoms and diagnoses of the fault diagnosis system. The rest of  
8 Figure 3 is the logic of the fault diagnosis system which resides in the local agent. The  
9 logic is identical to that of Figure 2 (and is therefore labeled identically as in Figure 2)  
10 with two exceptions. Function block 210 of Figure 3, which asks the repository to  
11 identify a root graph based on the initially described symptoms, now incorporates logic to  
12 ask it remotely at a known or discoverable location. Similarly, function block 220 of  
13 Figure 3, which retrieves a given graph for processing by the diagnosis logic, must now  
14 incorporate logic to retrieve it remotely from the same known or discoverable location.  
15 Once a graph is retrieved, it is processed normally according to the prior description of  
16 Figure 2.

17 An example embodiment of this invention uses Java classes to store the graphs of  
18 the repository, load them on demand, and download them when the repository is remote.  
19 The object oriented nature of Java supports the aggregation and encoding of commonly  
20 occurring sequences of question-answer pairs and related symptoms and diagnoses into  
21 one Java class. Java's ability to dynamically load additional object code at run-time as  
22 Java classes (a quality known as "late-binding") supports the incremental loading of  
23 additional graphs as needed. And when the repository is remote, Java's ability to  
24 remotely download a class (or package of classes) in one transmission minimizes network  
25 interactions needed to retrieve subsequent question-answer pairs.

26 In another embodiment, code generated by any other programming language is  
27 used, especially if the programming language is object oriented and supports late binding  
28 and on-demand downloading of code as does Java in the previously described  
29 embodiment.

30 In another embodiment, the repository is downloaded as needed onto a "small"

1 device (i.e., a device with very limited memory, e.g., a PDA or cell-phone). By way of  
2 example, this is useful for on-site diagnosis of systems in the field without requiring  
3 dispatched field engineers to bring complete manual sets or large repositories of  
4 diagnostic information. In addition to receiving the diagnostic information in small  
5 chunks, the information received is also current as of the time of the download.

6 In another embodiment, the repository is hosted by a service provider, supporting  
7 the diagnostic procedures of a plurality of customers in a self-service mode, downloading  
8 graphs as needed from the service provider's hosting repository (via the Internet, for  
9 example).

10 In similar embodiments, the customer performing the diagnosis is replaced by a  
11 customer support center performing the diagnosis on behalf of its own customers, by a  
12 field representative on-site with the failing system, or by the faulty system itself in a  
13 self-diagnostic mode. The latter usage enables autonomic behavior, especially when the  
14 faulty system can take action to apply the solution ("self-heal") identified by the fault  
15 diagnosis system.

16 In another embodiment, the repository is used in an "on-demand" environment  
17 with the service provider charging each customer for the amount of resources consumed  
18 during the diagnosis session. In this environment, the number of graphs downloaded  
19 on-demand from the repository is a metric of system usage. When the embodiment of the  
20 invention also incorporates one object-oriented class per graph (as mentioned above), the  
21 metric reduces to the number of classes downloaded on-demand.

22 Thus the methods and systems of the present invention are suitable for e-business  
23 on demand, which includes providing customer services.

24 In another embodiment, the system being diagnosed is a machine, possibly a  
25 complex machine (e.g., a computer processor or an automobile) or a complex component  
26 (e.g., an automobile engine), although it might be a simpler machine or component (e.g.,  
27 a refrigerator door that does not close completely).

28 In similar embodiments, the system being diagnosed is either a software system or  
29 a process, either complex (e.g., a computer operating system that has failed or a sales  
30 presentation that failed to result in a successful sale) or simple (e.g. a child's computer

1 game that does not install properly).

2 In another embodiment, the diagnostic system is being used in a proactive manner  
3 to prevent faults before they occur (e.g., helping a customer to choose, or training a  
4 salesperson to sell, an appropriate product from a line of products that can be tailored to  
5 multiple uses; helping or training someone to install or configure hardware or software;  
6 or, helping a word processor user to format a document in a particular way). In this  
7 embodiment, the diagnostic system is also known as a help system or a training system  
8 and is particularly useful for customer self-help and employee self-training. Coupled  
9 with other features of the invention already mentioned, these users can be remote from  
10 the repository and be using even small devices for the training or help. Downloading  
11 graphs on demand also assures the latest content and style of interaction deemed useful by  
12 the content provider or service provider of the diagnostic, help, or training system for the  
13 target audience.

14 Thus, the invention provides a method for diagnosing from a repository at least  
15 one fault in a system, the repository represented as a directed graph having one or more  
16 undivided directed subgraphs. In an example embodiment of a method for diagnosing  
17 comprising the steps of: receiving a first description of the at least one fault; employing  
18 the first description to identify a response from the repository; if the response is a  
19 diagnosis stopping, otherwise identifying at least one subgraph responsive to the first  
20 description; using the at least one subgraph in determining the diagnosis, stopping if the  
21 diagnosis results, otherwise forming a modified description based upon the at least one  
22 subgraph; and replacing the first description with the modified description and repeating  
23 the steps of receiving, employing, identifying and using until the diagnosis results.

24 In some embodiments of the method: the first description is completely covered  
25 by the diagnosis, further comprising implementing a solution based on the diagnosis;  
26 and/or the first description includes a set of symptoms describing the at least one fault;  
27 and/or wherein the modified description includes an additional set of symptoms identified  
28 for probing by the subgraph; and/or the first subgraph is identified by a method  
29 employing an index mapping descriptions to initial subgraphs, the trivial index simply  
30 mapping all descriptions to one subgraph; and/or the repository is remote and the

1 undivided subgraphs are downloaded from the repository to a local agent performing the  
2 diagnosis as needed; and/or the system is a machine, a software program, a process, or  
3 any combination of these; and/or each undivided subgraph is implemented as executable  
4 code; and/or the executable code is written in an object-oriented programming language,  
5 and/or a programming language that supports late binding, and/or a programming  
6 language that supports on-demand downloading of classes; and/or the executable code is  
7 written in Java; and/or the local agent is a machine; and/or the diagnostic system is being  
8 used in a proactive manner to prevent faults before they occur as in a help or training  
9 system; and/or the remote repository is downloaded as needed onto a “small” device (i.e.,  
10 a device with very limited memory); and/or the remote repository is hosted by a service  
11 provider supporting a plurality of customers with each customer downloading subgraphs  
12 as needed to perform diagnosis; and/or at least one of the customers is a customer support  
13 center diagnosing faulty systems on behalf of a plurality of its own customers; and/or the  
14 customer is a field representative performing diagnosis of a failing system; and/or the  
15 customer is a faulty system operating in a self-diagnostic mode; and/or the faulty system  
16 applies the solution identified by the fault diagnosis system in a “self-healing” mode.

17 In some embodiments, the invention is an article of manufacture comprising a  
18 computer usable medium having computer readable program code means embodied  
19 therein for causing diagnosis from a repository of at least one fault in a system, the  
20 computer readable program code means in the article of manufacture comprising  
21 computer readable program code means for causing a computer to effect the steps of  
22 method of the present invention.

23 In some embodiments, the invention is a program storage device readable by  
24 machine, tangibly embodying a program of instructions executable by the machine to  
25 perform method steps for diagnosing from a repository at least one fault in a system.

26 In some embodiments, the invention is an apparatus comprising means for  
27 diagnosing from a repository at least one fault in a system. The repository is represented  
28 as a directed graph having one or more undivided directed subgraphs, the means for  
29 diagnosing comprising: means for receiving a first description of the at least one fault;  
30 means for employing the first description to identify a response from the repository;

1 means for if the response is a diagnosis stopping, otherwise identifying at least one  
2 subgraph responsive to the first description; means for using the at least one subgraph in  
3 determining the diagnosis, stopping if the diagnosis results, otherwise forming a modified  
4 description based upon the at least one subgraph; and means for replacing the first  
5 description with the modified description and repeating the steps of receiving, employing,  
6 identifying and using until the diagnosis results.

7 While the invention has been described in terms of particular embodiments, those  
8 skilled in the art will recognize that the invention can be practiced with modification  
9 within the spirit and scope of the appended claims.

10 Variations described for the present invention can be realized in any combination  
11 desirable for each particular application. Thus particular limitations, and/or embodiment  
12 enhancements described herein, which may have particular advantages to a particular  
13 application need not be used for all applications. Also, not all limitations need be  
14 implemented in methods, systems and/or apparatus including one or more concepts of the  
15 present invention.

16 The present invention can be realized in hardware, software, or a combination of  
17 hardware and software. A visualization tool according to the present invention can be  
18 realized in a centralized fashion in one computer system, or in a distributed fashion where  
19 different elements are spread across several interconnected computer systems. Any kind  
20 of computer system - or other apparatus adapted for carrying out the methods and/or  
21 functions described herein - is suitable. A typical combination of hardware and software  
22 could be a general purpose computer system with a computer program that, when being  
23 loaded and executed, controls the computer system such that it carries out the methods  
24 described herein. The present invention can also be embedded in a computer program  
25 product, which comprises all the features enabling the implementation of the methods  
26 described herein, and which - when loaded in a computer system - is able to carry out  
27 these methods.

28 Computer program means or computer program in the present context include any  
29 expression, in any language, code or notation, of a set of instructions intended to cause a  
30 system having an information processing capability to perform a particular function



1 either directly or after conversion to another language, code or notation, and/or  
2 reproduction in a different material form.

3 Thus the invention includes an article of manufacture which comprises a  
4 computer usable medium having computer readable program code means embodied  
5 therein for causing a function described above. The computer readable program code  
6 means in the article of manufacture comprises computer readable program code means  
7 for causing a computer to effect the steps of a method of this invention. Similarly, the  
8 present invention may be implemented as a computer program product comprising a  
9 computer usable medium having computer readable program code means embodied  
10 therein for causing a a function described above. The computer readable program code  
11 means in the computer program product comprising computer readable program code  
12 means for causing a computer to effect one or more functions of this invention.

13 Furthermore, the present invention may be implemented as a program storage device  
14 readable by machine, tangibly embodying a program of instructions executable by the  
15 machine to perform method steps for causing one or more functions of this invention.

16 It is noted that the foregoing has outlined some of the more pertinent objects and  
17 embodiments of the present invention. This invention may be used for many  
18 applications. Thus, although the description is made for particular arrangements and  
19 methods, the intent and concept of the invention is suitable and applicable to other  
20 arrangements and applications. It will be clear to those skilled in the art that  
21 modifications to the disclosed embodiments can be effected without departing from the  
22 spirit and scope of the invention. The described embodiments ought to be construed to  
23 be merely illustrative of some of the more prominent features and applications of the  
24 invention. Other beneficial results can be realized by applying the disclosed invention in  
25 a different manner or modifying the invention in ways known to those familiar with the  
26 art.